Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 23 January 2014 (MN LATEX style file v2.2)

A spectral synthesis code for rapid modelling of supernovae

Wolfgang E. Kerzendorf[1], Stuart A. Sim[2]
[1] *Department of Astronomy and Astrophysics, University of Toronto, 50 Saint George Street, Toronto, ON M5S 3H4, Canada*
[2] *Astrophysics Research Centre, School of Mathematics and Physics, Queen's University Belfast, Belfast BT7 1NN, UK*

arXiv:1401.5469v1 [astro-ph.SR] 21 Jan 2014

23 January 2014

ABSTRACT

We present TARDIS - an open-source code for rapid spectral modelling of supernovae (SNe). Our goal is to develop a tool that is sufficiently fast to allow exploration of the complex parameter spaces of models for SN ejecta. This can be used to analyse the growing number of high-quality SN spectra being obtained by transient surveys. The code uses *Monte Carlo* methods to obtain a self-consistent description of the plasma state and to compute a synthetic spectrum. It has a modular design to facilitate the implementation of a range of physical approximations that can be compared to asses both accuracy and computational expediency. This will allow users to choose a level of sophistication appropriate for their application. Here, we describe the operation of the code and make comparisons with alternative radiative transfer codes of differing levels of complexity (SYN++, PYTHON, and ARTIS). We then explore the consequence of adopting simple prescriptions for the calculation of atomic excitation, focussing on four species of relevance to Type Ia supernova spectra – Si II, S II, Mg II, and Ca II. We also investigate the influence of three methods for treating line interactions on our synthetic spectra and the need for accurate radiative rate estimates in our scheme.

Key words: radiative transfer – methods: numerical – supernovae: general

1 INTRODUCTION

The goal of supernova (SN) modelling is typically to infer the composition, mass and explosion energy of an event with the aim of understanding the explosion mechanisms and progenitor systems. Studies often rely on radiative transfer / spectrum synthesis codes to interpret the complex spectra originating in rapidly expanding and often metal-rich ejecta. A variety of different codes exist and are used for different purposes: these range from very simple approaches designed purely for line identification [run times of second e.g. SYN++ (Thomas et al. 2011)] via computationally inexpensive one-dimensional codes [runtimes of minutes, e.g. the Mazzali & Lucy (1993, ML93) code] to vastly more complex codes that implement sophisticated micro-physics, time-dependence and/or multi-dimensionality (e.g. Pauldrach et al. 1996; Baron et al. 1996; Lentz et al. 2001; Höflich 2003; Dessart & Hillier 2005; Kasen et al. 2006; Kromer & Sim 2009; Blondin et al. 2013; Sauer et al. 2006). Such sophisticated approaches, however, come at the expense of execution times running to many hours (even days) on massively parallel machines, making them impractical for efficient exploration of parameter space or rapid characterization of new data sets. Thus, for the foreseeable future, the need for the simpler approaches will remain. Nevertheless, given the quality of modern data sets, it is important to test and understand the implications of any simplifying assumptions that are made.

Our goal is to provide analysis tools that strike an optimal balance between accuracy and computational expense, ideally in a manner that can be adjusted to suit the demands of (and resources available to) a variety of studies. To this end, we have undertaken the development of a modular SN radiative transfer code (TARDIS). As input, the code takes a one-dimensional model for the supernova eject (which can include arbitrary density and/or abundance stratification, if required) together with a luminosity and time since explosion. From these, ionization and excitation states are iteratively estimated and a synthetic spectrum is calculated. The code serves two main purposes. First, it is designed to have an execution time of minutes, which allows for rapid exploration of parameter space with the possibility of fitting and interpretation of observational data. Secondly, it serves as a platform to trial different approximations for important microphysics (e.g. line interaction, level populations, etc.). TARDIS allows for swift implementation and testing of new physics that – if successful – might be used in the more complete but complex codes like ARTIS (Kromer & Sim 2009) and SEDONA (Kasen et al. 2006), which are founded on the same principles as TARDIS. Even if the newly implemented physics is computationally too costly to use in large parameter-space studies, it can be still be used to quantify the uncertainty associated with using faster, but less accurate, approaches (i.e. Uncertainty Quantification).

In Section 2, we present the numerical methods and a subset of the modes of operation currently implemented in the code. Sources of atomic data are described in Section 3. Sections 4 and 5 present results of numerical convergence tests and detailed comparisons of our synthetic spectra with those from a selection of existing radiative transfer codes of varying degrees of complexity.

In this paper we explore the effect of approximations used by

two different physics sub-modules, focussing on their importance in the modelling of Type Ia supernovae (SNe Ia).

Studies of SNe Ia using 1D codes that combine realistic treatments of ionization and excitation (e.g. Baron et al. 1996; Pauldrach et al. 1996; Blondin et al. 2013) have already shown that the effects of departures from local thermodynamic equilibrium (LTE) can be rather important. Consequently, for our goal of developing a code that is both reasonably fast and accurate, we must attempt to identify suitable approaches and quantify the impact of approximations made in the interests of computational expediency. Thus, we first (in Section 6.1) examine the influence of adopting a simple analytic treatment for the populations of excited atomic/ionic levels by comparing to results using level populations obtained by numerically solving a set of equations of statistical equilibrium.

Secondly, in Section 6.2, we compare results obtained using different levels of sophistication in the treatment of line scattering and fluorescence. It has been clearly established that non-resonance line scattering has a key role in shaping the spectra of metal-rich SN ejecta (Mazzali 2000; Pinto & Eastman 2000; Baron et al. 1996; Dessart & Hillier 2005; Blondin et al. 2013). In this work, we will compare two versions of the "macro atom" scheme (Lucy 2005; Kromer & Sim 2009) to a resonance scattering scheme and quantify their influence on synthetic spectra.

In Section 7, we summarise our findings and outline future plans for the further development of this project.

2 METHOD

The TARDIS code is a new implementation of indivisible energy-packet *Monte Carlo* (MC) methods (Abbott & Lucy 1985; Lucy & Abbott 1993, ML93, Lucy 1999, 2002, 2003) that have previously been used to model radiation transport in Type Ia supernovae (SNe Ia; e.g. ML93; Kasen et al. 2006; Sim 2007; Kromer & Sim 2009). The algorithm operates by seeding a set of MC quanta representing photon bundles (r-packets in the nomenclature of Lucy 2002) and following their propagation through a model for the SN ejecta, as described below.

TARDIS is based on the same methods used by ARTIS, but the scope of TARDIS is much more limited so that runtimes are orders of magnitude shorter. In particular, having in mind the goal of rapid fitting of observations (i.e. using TARDIS to compute synthetic spectra that can be fed into a fitting algorithm and explore model parameter space), TARDIS currently neglects multi-dimensionality and time-dependence. Thus, a TARDIS calculation provides a single snapshot spectrum while an ARTIS run provides a time series of spectra from which synthetic light curves can be constructed.[1] In addition, TARDIS adopts an inner boundary approximation, which greatly simplifies the calculations by avoiding the need to describe the diffusion of radiation at very high optical depths.

TARDIS already includes a number of different options and modules that make it possible to compare the effects of particular assumptions on both the outputs and the computational expense. The current options will each be introduced in the sections below but, for reference, we provide a summary in Table 1. So far, the development of TARDIS has focused on SN Ia applications. However, we stress that our goal is to continue development and that the code has been structured to make it easy to implement additional physics and improve, alter or lift many of the existing approximations (see below). Further modules extending the code for application to Type II supernovae (SNe II) are being developed (Klauser, Kromer et al., in prep.).

TARDIS is written in Python with C extensions and is available as open source (BSD 3 clause license) at `http://pypi.python.org/pypi/tardis-sn`. For all random numbers TARDIS uses the Mersenne Twister random number generator (Matsumoto & Nishimura 1998). In Appendix A we show an example input file and provide basic operational details for using TARDIS. However, we refer potential users to the technical manual (available at `http://tardis.rtfd.org`) for further information on running the code.

2.1 Model Setup

In the interest of computational expedience, TARDIS approximates the SN ejecta as spherically symmetric such that all physical properties depend only on the radial coordinate r. The computational domain is defined by inner and outer radial boundaries. The code currently assumes that material inside the domain is in radiative equilibrium and neglects any non-radiative energy sources (i.e. it is assumed that energy injection by ^{56}Ni within the computation volume is negligible). This approximation limits the applicability of TARDIS to epochs at which the effective photosphere is external to the volume in which the majority of the luminosity is generated (for SNe Ia, this is most valid at early times and becomes an increasingly poor approximation at later epochs).

The computational domain is discretised into multiple cells (spherical shells). For each of these, the density and elemental abundances must be specified (the "input model"). The input model is setup at runtime and can involve any density profile and set of stratified abundances. Simple models with uniform abundances and standard density profiles are directly created by the code (see Appendix A) while models with stratified abundances or arbitrary density profiles are setup by providing simple input files, as described in the manual. During the simulation, various quantities are computed on a cell-by-cell basis, most importantly the parameters of a simple radiation-field model (see Section 2.2), which are used to estimate ionization fractions and level populations (see Section 2.3).

We assume that the ejecta are in homologous expansion, which becomes an adequate approximation for SNe Ia within $\sim$ 100 seconds after explosion (Röpke 2005).

2.2 Radiation Field Model

In general, we assume that both ionization and excitation are primarily controlled by the radiation field and we follow Abbott & Lucy (1985), Lucy & Abbott (1993) and ML93 in adopting approximations based on a simple model for the frequency (ν) dependent mean intensity

$$J_\nu = W B_\nu(T_\mathrm{R}) \tag{1}$$

where, B_ν is the Planck function. The radiation temperature (T_R) and dilution factor (W) are parameters that are iteratively derived from MC estimators in each grid cell (see Section 2.7).

[1] We note, however, that TARDIS can be used to compute a time series of synthetic spectra for a model via a sequence of calculations in which the input parameters specifying luminosity and time since explosion are varied.

Table 1. Summary of the TARDIS modes of operation used in this study.

Mode name	Mode of operation
Ionization modes:	
`lte`	Ionization fractions from the Saha-Boltzmann equation and $T_{\rm R}$ (Equation 2)
`nebular`	Ionization fractions from the nebular approximation (Equation 3)
Excitation modes:	
`lte`	Level populations from Boltzmann equation and $T_{\rm R}$ (Equation 4)
`dilute-lte`	Excited level populations reduced by W (Equation 5)
`nlte`	Population ratios determined from statistical equilibrium (Section 2.3)
Radiative rate modes:	
`dilute-blackbody`	Radiative rates obtained using $J^b_{lu} = W B_{\nu_{lu}}(T_{\rm R})$
`detailed`	Radiative rates obtained using MC estimators for J^b_{lu} (Section 2.7)
Line interaction modes:	
`scatter`	All line interactions treated as resonance scattering events
`downbranch`	Macro atom scheme with all internal transition probabilities set to zero (Section 2.5)
`macroatom`	Full macro atom scheme (Section 2.5)

2.3 Plasma State

To compute opacities and handle radiation-matter interactions, it is necessary to specify the number densities for the states of the atomic/ionic species included. In principal, these could be obtained by a full solution of the complete system of equations of statistical equilibrium (as done by Pauldrach et al. 1996; Baron et al. 1996; Blondin et al. 2013) with rate coefficients for radiative processes derived from properties of the MC radiation field (Lucy 2002, 2003). This will be implemented in future versions of TARDIS. However, since the computational cost of such an approach will be significant, it is also valuable to consider simpler approximations that can be used for rapid calculations to explore parameter space. We pursue the approximate approach here, closely following ML93 who have demonstrated the utility of such a philosophy (see also Pauldrach et al. 1996).

Currently, TARDIS determines the ionization balance for each of the model cells based on the density, elemental abundance, radiation temperature ($T_{\rm R}$), electron temperature ($T_{\rm e}$) and the dilution factor (W). In the calculations presented in this work, we will adopt either a standard Saha-Boltzmann equation (`lte` ionization mode),

$$\frac{N_{i,j+1} n_e}{N_{i,j}} = \Phi_{i,j,T_{\rm R}} \equiv \frac{2 Z_{i,j+1}(T_{\rm R})}{Z_{i,j}(T_{\rm R})} \left(\frac{2\pi m_e k T_{\rm R}}{h^2} \right)^{3/2} e^{-\frac{\chi_{i,j}}{k T_{\rm R}}} \tag{2}$$

or a modified nebular approximation (ML93), referred to as `nebular` ionization mode

$$\frac{N_{i,j+1} n_e}{N_{i,j}} = W[\delta\zeta_{i,j} + W(1-\zeta_{i,j})] \left(\frac{T_{\rm e}}{T_{\rm R}} \right)^{1/2} \Phi_{i,j,T_{\rm R}} \tag{3}$$

Here, n_e is the free electron number density, $N_{i,j}$ is the ion number density, $Z_{i,j}$ the partition function and $\chi_{i,j}$ the ionization potential for ion j of element i. $\zeta_{i,j}$ is the fraction of recombinations that go directly to the ground state in recombination to ion j of element i. δ is a correction factor introduced to approximately account for the dominance of locally created radiation at short wavelengths (see ML93). We follow ML93 in setting $T_e = 0.9 T_R$.

With ionization ratios determined using one of these approximations, the code solves for a complete set of ion number densities (and the electron density) by enforcing the appropriate total elemental number densities and insisting on charge conservation in the usual manner.

TARDIS can use a variety of approximations for calculating the level populations within each ion. As a simplest case, we adopt the Boltzmann excitation formula,

$$n_{i,j,k} = \frac{g_{i,j,k}}{Z_{i,j}} N_{i,j} \exp(-\epsilon_{i,j,k}/k T_{\rm R}), \tag{4}$$

where $n_{i,j,k}$ is the number density of level k of ionization state j of element i, which has excitation energy $\epsilon_{i,j,k}$ relative to the ion ground state. We will refer to this as `lte` excitation mode.

As a slightly more sophisticated approach, TARDIS includes a crude NLTE approximation for level populations in which we continue to apply Equation 4 to all metastable levels but adopt

$$n_{i,j,k} = W \frac{g_{i,j,k}}{Z_{i,j}} N_{i,j} \exp\left(-\epsilon_{i,j,k}/k T_{\rm R}\right) \tag{5}$$

for all other excited state (see Equation 4 in Lucy 1999). We refer to this approach as the `dilute-lte` excitation mode in TARDIS.

Finally, TARDIS is also able to obtain NLTE level populations by solving a set of statistical equilibrium equations for chosen species (`nlte` excitation mode). Here we formulate rates between pairs of levels (upper level u and lower level l)[2] in an ion as

$$\mathcal{R}_{ul} = \beta_{lu} A_{ul} n_l + \beta_{lu} B_{ul} n_l J^b_{lu} + C_{ul} n_u n_e \tag{6}$$

$$\mathcal{R}_{lu} = \beta_{lu} B_{lu} n_l J^b_{lu} + C_{lu} n_l n_e, \tag{7}$$

where A, B are the usual Einstein coefficients for radiative transitions, C is the rate coefficient for electron collisions and $\beta_{lu} = \frac{1}{\tau_{lu}}[1 - \exp(-\tau_{lu})]$ is the Sobolev escape probability (see Lucy 2002, and Section 2.5). J^b_{lu} is the mean intensity at the extreme violet wing of the bound-bound transition between levels l and u (see Lucy 2003). TARDIS can currently estimate J^b_{lu} in one of two ways. The first option (`dilute-blackbody`) sets $J^b_{lu} = W B_{\nu_{lu}}(T_{\rm R})$ (i.e. imposes the radiation-field model of Section 2.2). Alternatively, we have also implemented a `detailed` setting that calculated J^b_{lu}-values using estimators (see Section 2.7).

The net rate of change for a level population is

$$\frac{dn_k}{dt} = \sum_{i \neq k} R_{ik} - \sum_{i \neq k} R_{ki} \ . \tag{8}$$

[2] For clarity, we drop the explicit reference to atomic number and ionization number here.

In `nlte` excitation mode, TARDIS solves for level populations satisfying $\frac{dn_k}{dt} = 0$ for all k (with $\sum_k n_k = N$, the total ion population). In practise, the code is initialized with LTE level populations and initial guesses for β_{lu} and J^b_{lu}. It then solves the statistical equilibrium equations using LU decomposition (LAPACK; Anderson et al. 1999). After each MC iteration (see Section 2.8) the statistical equilibrium equations are re-solved using the last level populations to compute values for β_{lu} and results from the MC step to estimate J^b_{lu}. This process eventually converges to a stable set of level populations (see Section 2.8).

Note that our current implementation includes only bound-bound rates in the set of statistical equilibrium equations – bound-free processes are neglected (but will be implemented as part of future ionization modules). Consequently, our `nlte` excitation mode will not yield accurate populations for levels that are significantly (de)-populated via bound-free processes. Nevertheless, our `nlte` excitation mode is certainly an improvement over LTE (or dilute-LTE) populations and provides a convenient means by which to quantify the plausible scale of errors introduced by the simple analytic formulae (Section 6.1).

2.4 Initialization of Monte Carlo quanta

For each MC simulation of the radiation field, we begin by initializing a population of r-packets at the inner boundary of the computational domain. Each r-packet is an indivisible quantum of radiative energy with an associated photon frequency (ν). We choose to initialize all our r-packets with equal co-moving frame energy E.

Currently it is assumed that the radiation field injected through the inner boundary has a black body frequency distribution at temperature $T_{\rm i}$ such that the luminosity launched at the inner boundary is

$$L_i = 4\pi r_{\rm i}^2 \sigma T_{\rm i}^4 \tag{9}$$

where $r_{\rm i}$ is the radius of the inner boundary. Consequently, the co-moving frame frequencies of the packets are selected by randomly sampling the Planck function (for $T_{\rm i}$) and the packet energies are chosen to match L_i: i.e., in a simulation with N r-packets, the co-moving frame energy of each packet is initialised to

$$E = \frac{4\pi r_{\rm i}^2 \sigma T_{\rm i}^4}{N} \Delta t \tag{10}$$

where Δt is the (arbitrarily chosen) time interval represented by the simulation. The procedure adopted to choose $T_{\rm i}$ is described in Section 2.8. The r-packets are assigned initial directions of propagation specified by $\mu = \sqrt{z}, z \in (0, 1]$, as appropriate for zero limb darkening. [3]

2.5 Radiation-matter interactions

In our indivisible energy packet scheme, opacity does not destroy MC quanta but can change both the photon frequency and direction of propagation associated with a packet. The co-moving frame energy of the packet is always conserved during interaction. This ensures that radiative equilibrium is strictly enforced throughout the simulation.

[3] We specify direction of propagation by $\mu = \cos\theta$, where θ is the angle between the packet flight path and the radial direction.

Currently, TARDIS treats only two classes of radiation-matter interaction – Thomson scattering by free electrons and bound-bound interactions with atoms/ions. For Thomson scattering the opacity encountered by an r-packet in a path length s in direction μ is given by

$$\tau_{\rm e} = \kappa_e s \tag{11}$$

where $\kappa_e = \sigma_{\rm T} n_{\rm e} D_\mu$ is the observer frame opacity coefficient. Here, $\sigma_{\rm T}$ is the Thomson cross section and $n_{\rm e}$ is the number density of free electrons. We include the first order Doppler factor $D_\mu = (1 - \mu v/c)$ to preserve $O(v/c)$ accuracy in transforming between observer and co-moving frames (following Lucy 2005). TARDIS treats Thomson scattering as a coherent scattering process: it causes packets to change their direction but not their co-moving frame frequency or energy. Following a Thomson scattering event, the new direction is drawn from an isotropic distribution, $\mu = -1 + 2z, z \in [0, 1]$.

Bound-bound transitions are the dominant (and also most complex) opacity source in SN Ia ejecta (Pinto & Eastman 2000). TARDIS treats bound-bound opacity in the Sobolev approximation (see e.g. Lamers & Cassinelli 1999), which is appropriate for media with large velocity gradients. In the case of homologous expansion, the Sobolev optical depth of the transition between lower state l and upper state u is given by

$$\tau_{\rm S} = \frac{\pi e^2}{m_e c} f \lambda_{lu} t_{\rm exp} n_l \left(1 - \frac{g_l n_u}{g_u n_l}\right), \tag{12}$$

where f is the absorption oscillator strength of the transition, λ_{lu} is the wavelength and $t_{\rm exp}$ the time since explosion. The last term (in brackets) corrects for stimulated emission.

During the MC simulations (see Section 2.6), r-packets can be absorbed by bound-bound transitions. Currently, three approaches are implements in TARDIS to described the re-emission of an absorbed r-packet. In the simplest approach (mode `scatter`), all bound-bound interactions are treated as resonance scattering events [cf. Branch et al. 1985; Lucy & Abbott 1993; Long & Knigge (2002, LK02)]. Consequently, absorbed r-packets are simply re-emitted with the same co-moving frame frequency but a new direction of propagation (determined by randomly sampling an isotropic distribution, as above).

Our most sophisticated approach to line interactions is an implementation of the macro atom scheme (mode `macroatom`) devised by Lucy (2002, 2003) and used by Kromer & Sim (2009). In the `macroatom` mode, bound-bound absorption of an r-packet activates a macro atom to the upper level of the absorbing transition. The macro atom algorithm is then used to simulate the re-emission of the absorbed energy, in accordance with the assumptions of radiative and statistical equilibrium. Currently, our macro atom implementation includes only bound-bound radiative transitions although we stress that the method can be readily extended to account for collisional and/or bound-free processes if suitable atomic data are provided. When a macro atom is activated, it can undergo a sequence of internal state transitions until ultimately deactivating and returning an r-packet to the MC simulation (see Lucy 2002, 2003).

Our third method of handling bound-bound interaction is a simplified version of macro atom, referred to as `downbranch`. This version is effectively the same as the `macroatom` scheme except that the probabilities for all internal transitions are set to zero. Thus macro atoms are forced to deactivate directly from the

state originally activated. This approach avoids some of the computational cost of the full macro atom machinery while retaining a simplified treatment of fluorescent processes, akin to that introduced by Lucy (1999).

2.6 Propagation of Monte Carlo quanta

In the current implementation, there are three processes that can act to terminate the flight path of an r-packet: reaching a grid zone boundary, undergoing electron scattering or being absorbed by a bound-bound transition. Choosing which class of event occurs is determined via a simple MC experiment. First, an optical depth ($\tau_{\rm n}$) that the r-packet may propagate without being absorbed is randomly selected in accordance with the $\exp(-\tau_{\rm n})$ attenuation law: $\tau_{\rm n} = \ln z, z \in (0, 1]$. Next we identify the closest redward line transition to the co-moving frame frequency of the r-packet. Then we calculate the distance ($d_{\rm l}$) the packet must travel to Doppler-shift into Sobolev resonance with that line. We also calculate $d_{\rm e}$, defined by $\tau_{\rm n} = \kappa_e d_{\rm e}$. Finally, we calculate the distance $d_{\rm s}$ the r-packet would need to propagate to reach the boundary of the current shell.

To determine which class of event will terminate the flight path, we compare the three distances $d_{\rm l}$, $d_{\rm e}$ and $d_{\rm s}$. If $d_{\rm l}$ is the shortest, then bound-bound absorption is possible. To test whether line absorption occurs, we compare $\tau_{\rm tot} = \tau_{\rm S} + \sigma_{\rm T} n_{\rm e} d_{\rm l}$ to $\tau_{\rm n}$. If $\tau_{\rm tot} > \tau_{\rm n}$ an interaction with the line will occur: we propagate the r-packet to the position of Sobolev resonance and then process the line interaction event in accordance with the procedures described in Section 2.5 for bound-bound absorption. Alternatively, if $\tau_{\rm tot} < \tau_{\rm n}$, we reduce $\tau_{\rm n} \rightarrow \tau_{\rm n} - \tau_{\rm S}$ and recompute $d_{\rm e}$. We also recompute $d_{\rm l}$, now considering the next redward line transition. With new values of $d_{\rm e}$ and $d_{\rm l}$ we again compare the three distances ($d_{\rm l}$, $d_{\rm e}$, $d_{\rm s}$) and continue this process until either a line absorption event is triggered or until $d_{\rm l}$ is no longer the shortest distance.

If $d_{\rm e}$ is shortest, an electron scattering event terminates the r-packet flight path. The packet is propagated through distance $d_{\rm e}$ and sent to the algorithm dealing with electron scattering events (see Section 2.5).

Finally, if $d_{\rm s}$ is shortest, the r-packet is propagated to the appropriate grid cell boundary and flagged as having successfully crossed that boundary. If the cell boundary is internal to the computational domain, the propagation of the packet through the new shell continues as before (noting that all material properties relevant to computing the opacity may have changed). When an r-packet reaches either of the boundaries of the computational domain (outer or inner), its flight path terminates and its final properties are recorded. r-packets reaching the outer boundary are assumed to the escape freely (such that they can contribute to the observable spectrum). Packets crossing the inner boundary are assumed to be reabsorbed by the inner ejecta and thus lost.

2.7 Monte Carlo estimators

During MC simulations, TARDIS uses the r-packet trajectories to collect estimators for radiation field properties, as required for calculation of the ionization and excitation conditions. Specifically, we record a pair of estimators (following ML93; LK02),

$$J_{\rm est} = \frac{1}{4\pi \Delta t V} \sum E l D_\mu \tag{13}$$

and

$$\bar{\nu} = \frac{1}{4\pi \Delta t V} \sum E \nu l D_\mu \tag{14}$$

for each grid cell. The summation is made over all r-packet trajectories inside the cell. For each trajectory, E and ν are the packet energy and frequency (in the co-moving frame) and l is the length of the trajectory (measured in the observer frame). After each MC calculation, we use these estimators to obtain new values for the parameters of our radiation field model (see Section 2.2):

$$T_{\rm R} = \frac{h}{k_B} \frac{\pi^4}{360\zeta(5)} \frac{\bar{\nu}}{J_{\rm est}} \tag{15}$$

and

$$W = \frac{\pi J_{\rm est}}{\sigma_{SB} T_{\rm R}^4} \ . \tag{16}$$

In `detailed` mode, we also record estimators for the mean intensity at the violet wing of each line transition J^b_{lu} in each cell (Lucy 1999):

$$J^b_{lu} = \frac{1}{4\pi \Delta t V} \frac{t_{exp}}{c} \sum \frac{E}{\nu} D_\mu \tag{17}$$

where the summation is now over all r-packets that pass through Sobolev resonance with the $l \rightarrow u$ transition in the cell.

2.8 Iteration cycle

TARDIS performs a sequence of MC simulation during which the values of $T_{\rm i}$, $T_{\rm R}$ and W (and, in `detailed` mode, J^b_{lu}) are iteratively improved until convergence is reached.

To initialize a calculation, we set $T_{\rm R}$ (in all shells) to a sensible guess, typically 10000 K. The dilution factor is initialized to follow geometric dilution, $W = \frac{1}{2}[1 - (1 - (r_{\rm i}/r)^2)^{1/2}]$, where r is the radius at the centre of each shell and $r_{\rm i} = v_{\rm i} t_{exp}$ is the radius of the inner boundary. In `detailed` mode, we initialize $J^b_{lu} = W B_{\nu_{lu}}(T_{\rm R})$. Although $T_{\rm i}$ can also be treated as a simple input parameter, we anticipate the goal of fitting observations and so generally adopt the luminosity at the *outer* boundary (L_o) as a simulation parameter. As a first guess we then adopt

$$T_{\rm i} = \left(\frac{L_{\rm o}}{4\pi r_{\rm i}^2 \sigma} \right)^{1/4} \ . \tag{18}$$

A MC simulation is then performed (see Section 2.6) and the resulting estimators (see Section 2.7) are used to update the plasma properties (see Section 2.3) in each shell. We compare the total energy emerging through the outer boundary to the requested value of $L_{\rm o}$ and modify $T_{\rm i}$ (and therefore the luminosity at the inner boundary) to obtain better agreement (i.e., if the emergent luminosity is lower than requested, $T_{\rm i}$ is increased and vice versa). Using these updated quantities, a new MC calculation is carried out and used to update the model properties again. This process is repeated until the plasma state has converged to sufficient accuracy (or a chosen maximum number of iterations is reached); convergence properties are described for example calculations in Section 4.

2.9 Synthetic spectrum

Synthetic spectra can be obtained directly from MC radiative transfer calculations by binning the quanta that emerge through the outer boundary in frequency-space. However, Lucy (1999) and Mazzali (2000), showed that higher-quality synthetic spectra can be extracted from the simulations using only slightly more sophisticated algorithms and various such approaches have been used in MC codes (ML93; LK02; Sim & Mazzali 2008).

In TARDIS, we have implemented a simple approach to compute the synthetic spectrum, very similar to that adopted by LK02 and Sim et al. (2010). This method is fully consistent with the macro atom scheme (Lucy 2002, 2003). The spectrum is calculated during one final MC simulation that is performed once the iterative sequence of simulations discussed in Section 2.8 has converged. During this last simulation, whenever an r-packet is about to be launched on a new trajectory (either because it is about to be injected through the inner boundary, or because an interaction occurred inside the computational domain), the main simulation is suspended and a set of N_v test packets, which we refer to as *virtual* (v-) packets, are created. These v-packets have properties identical to those of the r-packet that was about to be launched, except that they each have a different propagation direction and their energies are assigned based on the probability distribution associated with the creation of the original r-packet (see below). Since the only purpose of v-packets is to estimate a contribution to the emergent spectrum, we do not spawn v-packets on trajectories that ultimately intersect the inner boundary. Thus, for v-packets created at radius r, we assign directions (μ_v) in the interval $\mu_{min} < \mu_v < 1$, where $\mu_{min} = -\sqrt{(1-(r_{\rm i}/r)^2)}$. For v-packets generated following a physical interaction inside the domain, the co-moving frame energy assigned (E_v) is simply

$$E_v = E\frac{1-\mu_{min}}{2N_v} \tag{19}$$

where E is the co-moving frame energy of the parent r-packet. For v-packets created when an r-packet is about to be injected through the inner boundary, the energies assigned are modified to account for the adopted angular distribution of the incoming radiation field

$$E_v = E\frac{2\mu_v}{N_v} \tag{20}$$

Once created, each v-packet is propagated though the simulation domain. Unlike r-packets, v-packets do not undergo physical interactions with the medium (they are never scattered), but the total optical depth along the trajectory of the$v$-packet (τ_v) is recorded. Each v-packet makes a contribution to the photon frequency-bin of the synthetic luminosity (ergs $\rm s^{-1}$ $\rm Hz^{-1}$) spectrum in which its observer frame frequency lies

$$L_v = \frac{E_v}{\Delta t \Delta\nu}\exp(-\tau_v) \tag{21}$$

where $\Delta\nu$ is the width of the spectral frequency bin. The spectrum computed from the v-packets agrees with that obtained by direct binning of emerging r-packets but typically has considerably less MC noise. The reduction in MC noise and the computational overhead introduced by the v-packet step are determined by the choice of N_v. For the calculations in this paper we typically use $N_v = 3$ which reduced the MC noise by a factor of ≈ 3.

3 DATASOURCES

Our radiative transfer calculations require input atomic data for the calculation of the ionization/excitation state of the plasma (ionization potentials and atomic models), bound-bound opacities and radiative rate coefficients (lines lists and oscillator strengths) and, in our `nlte` excitation mode, electron collision rate coefficients (thermally averaged collision strengths).

For the simulations presented below we use an atomic database that includes all elements with $Z \leq 30$. Atomic masses were taken from Wieser & Coplen (2011) and ionization threshold energies from NIST (Kramida et al. 2012). The modified nebular approximation (see Section 2.3 equation 3) requires that the fraction of recombinations to the ground state (ζ-factor) be specified. For this, values were extracted from the data compiled by LK02 for a fixed set of temperatures (2000 – 50000 K; 2000 K grid spacing). During code execution, linear interpolation between these values is used to obtain ζ as a function of temperature.

We have drawn atomic/ionic energy levels and bound-bound (line) transition data from two sources and constructed two atomic databases. In the first case, we populate the lines and levels database from the Kurucz line list (Kurucz & Bell 1995) omitting levels above the ionization threshold and lines with an oscillator strength $\log gf < -3$. This dataset contains no electron collision rate data. For ions not present in the Kurucz line list, we use atomic models consisting of only the ground state (statistical weights from NIST) and no line transitions. We use this dataset for several of the test calculations presented in Section 5.

Secondly, we constructed a data set in which the Kurucz line list data are replaced with data from CHIANTI 7.1 (Dere et al. 1997; Landi et al. 2012) for select species of relevance to SNe Ia: Si II, S II Mg II and Ca II. CHIANTI provides atomic models and data for both radiative and electron collision bound-bound transitions. For ions taken from CHIANTI, no cuts were applied to $\log gf$ in selecting radiative data. Electron collision rates were calculated using the CHIANTI atomic dataset for temperatures between 2000 K and 48000 K (in steps of 2000 K) and interpolated during code execution. This data set is used for most of the the calculations presented in Sections 4, 5 and 6.

The two atomic data sets used here have been constructed with the goals of this study in mind (code testing and simple differential tests of modelling assumptions). However, TARDIS is capable of handling much larger atomic datasets and can easily be extended to include additional atomic processes.

4 CONVERGENCE TESTS

As described in Section 2.8, TARDIS performs an iterative sequence of simulations. Here we describe tests verifying the convergence properties of this iteration procedure. For these tests, we have chosen to use a model that is based on the output of a 1D SN Ia explosion simulation – this allows us to study the operation of the code applied in the most complex regime available (model with arbitrary density profile and stratified abundances including a large number of elements). [4] Specifically, we adopt the model for a detonation

[4] We note, however, that TARDIS does not require such a complex input model (see Section 2 and Appendix A). In many applications, simpler models will be more efficient for exploration of parameter space and attempting to fit observations.

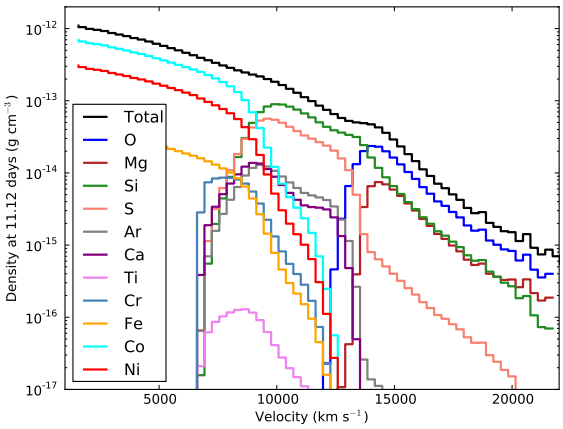

Figure 1. Densities for the 1.06 $M_\odot$ CO WD detonation model of Sim et al. (2010) at 11.1 days after explosion. The black line shows the total mass density and the coloured lines show the densities of important elements in the ejecta.

of a 1.06 $M_\odot$ white dwarf (WD) (Sim et al. 2010) that was computed for a uniform initial composition of ^{12}C and ^{16}O (50:50 mix with no ^{22}Ne). The density and composition for this model during the homologous phase are illustrated (at $t = 11.1$ days) in Fig. 1. The model has strongly layered ejecta, with intermediate mass elements (IMEs) dominating the composition above ~ 9000 km s^{-1} and iron-group elements dominating below. The complete model from Sim et al. (2010) consists of 70 shells. However, since inner and outer boundaries are imposed by TARDIS, typically only around half of the shells are included in the TARDIS calculations. For the test calculations we adopt $t_{exp} = 11.1$ days, $v_i = 11000$ km s^{-1}, $v_o = 22000$ km s^{-1} and $\log_{10} L_o/L_\odot = 9.34$ (the luminosity obtained at this epoch from ARTIS simulations – see Section 5.4).

We have tested convergence by carrying out a sequence of 30 iterations with 2×10^5 MC quanta per iteration. To test all aspects of the current implementation, we show results from calculations with the most sophisticated set of modes of operation: `macroatom` mode for line interactions, `detailed` mode for calculation of radiative rates and `nebular` ionization mode. We use `nlte` excitation mode for Si II and `dilute-lte` excitation for all other ions.

As described in Section 2.8, the code iteratively tries to match the requested SN luminosity (L_o) by changing the temperature of the black-body spectrum emitted from the inner boundary. This in turn changes the radiation field in the cells above (parametrized by T_R and W), which subsequently influences the ionization and excitation of the plasma. To speed convergence, we found it useful to adopt a scheme in which T_i is only changed after every third iteration: this allows T_R and W to respond to changes in T_i, helping the code make a more informed choice when modifying T_i in an attempt to match L_o.

Fig. 3 (upper panels) shows the rapid and uniform convergence for both T_R and W for our test calculation. By iteration 20 both quantities have ceased to evolve with iteration number. This is driven by the quick convergence of T_i (see Fig. 4).

We have tested the convergence properties of level populations for species treated with our `nlte` excitation mode by examining the evolution with iteration number of departure coefficients defined by

$$b_{\rm LTE} = \frac{n_i/n_0}{n_i^{\rm LTE}/n_0^{\rm LTE}}, \quad (22)$$

where n_i is the level population of the selected level and n_0 is that of the species ground state. This is illustrated for the $3s^2 4s\ ^2S_{1/2}$ state of Si II (see Fig. 9), in Fig. 3. Since the `nlte` excitation level populations are strongly affected by the radiative rates, we also show the convergence of J^b_{lu} for Si II λ6347 in Fig. 3. Although considerably more affected by MC noise than T_R or W, both J^b_{lu} and the departure coefficient have converged to a stable pattern after roughly 20 iterations. We explore the influence of MC noise in the estimators in Section 6.3.

We have tested the sensitivity of our results to the number of MC quanta by repeating the calculation described above using 10^6 and 2×10^6 MC quanta. The convergence of these runs was essentially identical to those of our calculation with 2×10^5 quanta and led to no changes in the output spectra (see Fig. 4). In addition, we show that the spectra extracted using *virtual* packets (see Section 2.9) are much less noisy than those obtained by directly extracting a spectrum by binning emergent quanta.

Finally, to test the influence of grid resolution, we carried out additional runs in which we increased the number of shells used by factors of two, five and ten (using 2×10^5 MC quanta in each case). We interpolated the velocities linearly (the entire model follows a strictly homologous velocity law), but did not interpolate the densities or abundances (we aim to test the influence of resolution on our computed radiation field properties, not the input model). Despite the finer sampling, the converged properties and the emergent spectrum for these tests were unaffected indicating that the resolution adopted in our standard calculation was adequate.

5 CODE COMPARISONS

We have made comparisons between TARDIS calculations and those of several other codes. We first focus on the calculation of the emergent spectrum (for fixed plasma conditions) by comparison with SYN++ (Section 5.1). We then use the CHIANTIPY 0.5.2 [5] package to test our solver for NLTE level populations in Section 5.2.

In Sections 5.3 and 5.4 we make more sophisticated tests by comparing TARDIS synthetic spectra to those of two alternative radiative transfer codes, PYTHON and ARTIS. Like TARDIS, these codes use MC methods and a Sobolev treatment of line opacity. However, they have different treatments of NLTE ionization, bound-free opacity and radiation sources (as detailed below). Comparing their results allows us both to identify potential issues (with any of the codes) and to quantify the effects of the different assumptions made by the codes.

5.1 Comparison to SYN++

Similar to TARDIS, SYN++ adopts a spherically symmetric SN model in homologous expansion. SYN++ handles line opacity in the Sobolev approximation and assumes all lines can be treated in the resonance scattering limit. In contrast to TARDIS, SYN++

[5] This work uses the version 0.5.2 with some corrections applied to the excitation calculation in collaboration with Ken Dere.

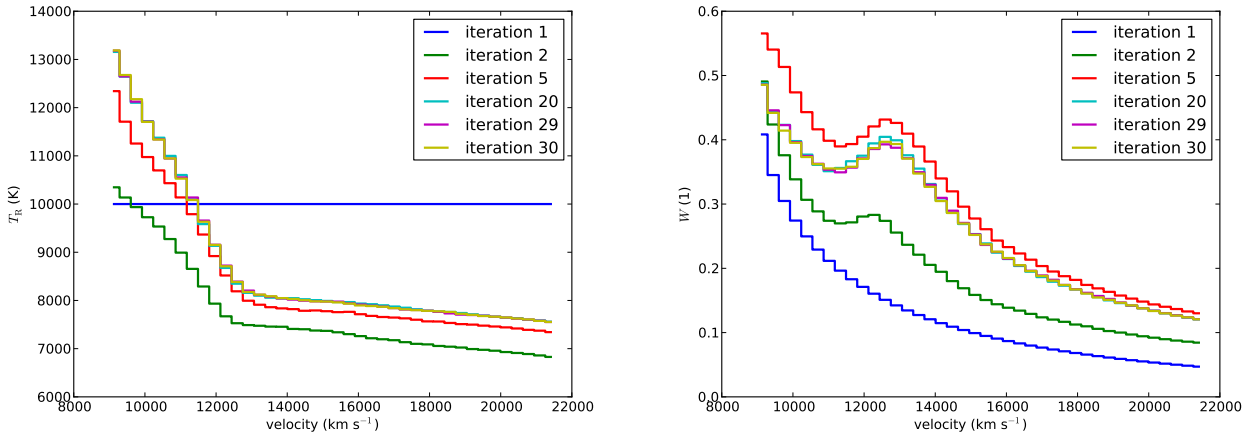

Figure 2. Convergence tests of T_{R} (left panel) and W (right panel) for a TARDIS calculation with 2×10^5 MC quanta.

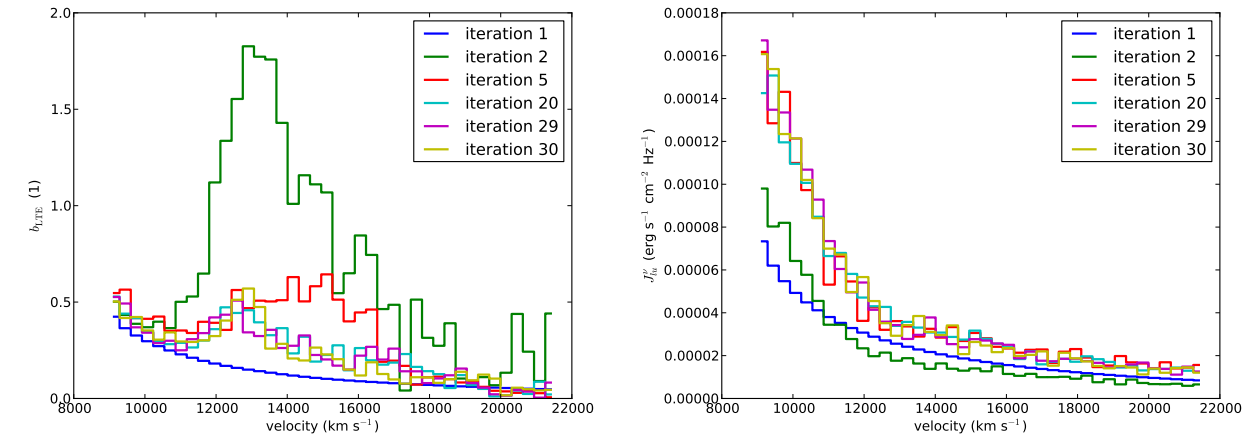

Figure 3. The left panel shows a convergence test for the departure coefficient b_{LTE} (see Section 5.2) of Si II $3s^2 4s\ ^2S_{1/2}$. For selected transitions, the rightpanel shows the convergence of J^b_{lu} calculated using MC estimators. The calculation used 2×10^5 MC quanta.

does not accept density and abundances as input parameters, but rather the opacity for a reference line associated with each ion species. SYN++ then calculates the opacities of other lines using the Boltzmann excitation formula at a user-specified temperature. Thus, comparison with SYN++ does not test the calculation of the plasma state (ionization/excitation) in TARDIS. However, it does provide a direct test of the accuracy to which the spectrum is computed for fixed plasma conditions.

For the test, we setup TARDIS with a pure silicon one-zone model. Line interactions were treated in `scatter` mode and electron scattering was disabled for ease of comparison with SYN++. Ionization and excitation were both treated in `lte` modes and we used the atomic data set drawn only from the Kurucz line list. The temperature of the inner boundary and of the temperature of the radiation field inside the only cell were set to 10000 K and 10^6 MC quanta were used to generate the TARDIS spectrum. The output τ_{S} for the reference lines of species Si I-IV were taken from TARDIS and used to calculate the spectrum with SYN++. In Fig. 5, we illustrate the excellent agreement between both codes, validating the method of spectrum formation used by TARDIS.

5.2 Testing of NLTE level populations

To test the accuracy of our NLTE level population scheme, we compare our results to those obtained using CHIANTIPY. For this test, we calculate the NLTE level populations for a plasma element of fixed electron density n_{e} and kinetic temperature T_{e} irradiated by a diluted black body ($W = 0.5$) with radiation temperature T_{R}. Note that, as in TARDIS, the current version of CHIANTIPY includes only bound-bound (radiative and electron collision) rates when solving for level populations (i.e. the influence of photo-ionization and recombination are not included).

In general we found excellent agreement between the calculation performed by TARDIS and CHIANTIPY. Fig. 6 shows the departure coefficient b_{LTE} for Si II computed for $T_{\mathrm{e}} = 9000$ K, $T_{\mathrm{R}} = 10000$ K and $n_{\mathrm{e}} = 4.15 \times 10^8 \mathrm{cm}^{-3}$. We find a bias of 0.006 and a standard deviation of 0.003 in the difference between the sets of departure coefficients obtained in the two calculations.

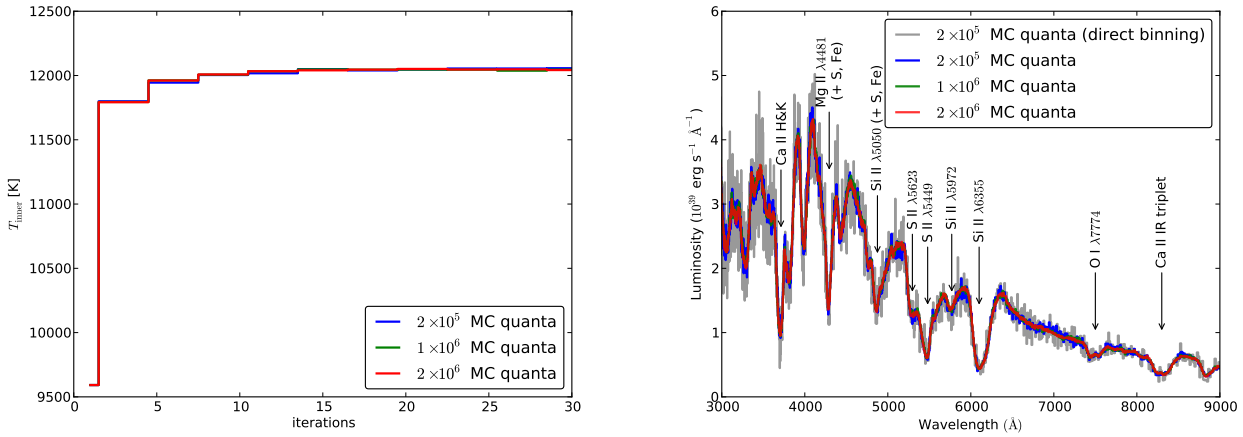

Figure 4. Results of $T_{\rm i}$ convergence tests with different numbers of MC quanta are shown on the left. The final spectra are shown on the right. For the 2×10^5 calculation we show both the spectrum obtained by directly binning the emergent MC quanta and that obtained using virtual packets (described in Section 2.9). For all other cases we show only the virtual packet spectrum. The positions of several prominent features (that will be discussed in later sections) are marked (and major blends are indicated, where appropriate).

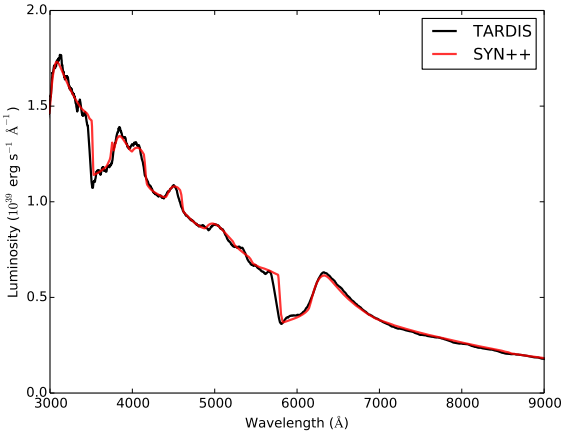

Figure 5. Comparison of SYN++ and TARDIS synthetic spectra for single cell pure silicon atmospheres. For ease of comparison we have convolved the TARDIS output using a Savitzky-Golay filter (Savitzky & Golay 1964).

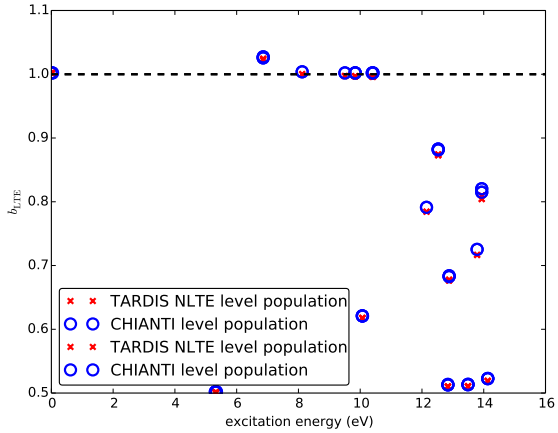

Figure 6. Comparison of Si II level populations calculated with CHIANTI and TARDIS as described in Section 5.2.

5.3 Comparison to PYTHON

PYTHON is a 2D MC radiative transfer code, which has been used to model synthetic spectra for accretion disk winds in a variety of astrophysical systems: cataclysmic variables (LK02; Noebauer et al. 2012), massive young stellar objects (Sim et al. 2005) and active galactic nuclei (Higginbottom et al. 2013). As in TARDIS, an input radiation field is specified and synthetic spectra are computed by tracking the propagation of radiation through a computational domain. PYTHON includes bound-bound, bound-free, free-free and electron scattering opacities. Although a `macroatom` treatment is implemented in PYTHON, to date this has only been used for modelling H I recombination lines (Sim et al. 2005): for metal lines, a two-level atom is used.

For our comparison calculation, we adopt a simple SN Ia ejecta model, following Lucy (1999). Specifically, we choose a density profile based on a simple fit ($\rho \propto v^{-7}$; cf. Branch et al. 1985) to the W7 model of Nomoto et al. (1984). We carry out the comparison for an epoch of 13 days post explosion, choosing an emergent bolometric luminosity of $\log_{10} L_{\rm o}/L_{\odot} = 9.44$ and setting the inner boundary of the computational domain at $v_{\rm i} = 11000$ km s^{-1}. We set the outer boundary at $v_{\rm o} = 20000$ km s^{-1}. Since PYTHON cannot currently handle position-dependent compositions, we adopt a uniform composition for the ejecta. At these epochs, the observed optical spectra of normal SNe Ia are dominated by features associated with IMEs. Therefore, we adopted an IME-rich composition (see Table 2) with relative abundances based on those found in modern SN Ia explosion models (specifically Seitenzahl et al. 2013). The TARDIS calculation was carried out using `nebular` ionization and `dilute-lte` excitation modes (see Section 2.3). For ease of comparison, we used our `scatter` mode for line opacity, which is our closest equivalent to the two-level treatment of metal lines used in PYTHON, and our atomic data set drawn from Kurucz line list (see Section 3).

For the PYTHON comparison run we set up the same model, using spherical geometry and imposing an homologous velocity law. The input radiation field for the PYTHON calculation was

Table 2. Elemental mass fractions adopted for the test model used in Section 5.3.

Element	Mass fraction	Element	Mass fraction
O	0.19	Mg	0.03
Si	0.52	S	0.19
Ar	0.04	Ca	0.03

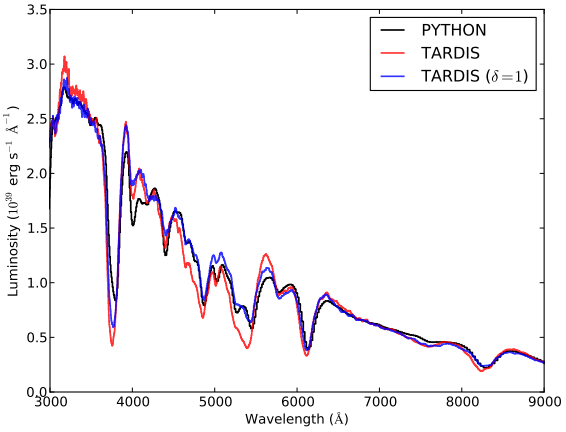

Figure 7. Synthetic spectra computed with PYTHON and TARDIS for the test model described in Section 5.3. The TARDIS calculations were carried out using our `scatter` and `nebular` ionization modes. In addition to our standard case, we also plot the spectrum obtained if we enforce $\delta = 1$ in TARDIS (see Equation 3).

forced to be the same as in the TARDIS calculation (i.e. we inject a black-body radiation field through the inner boundary with temperature set to be the same as that used by TARDIS). The PYTHON calculations were also carried out using an atomic data from the Kurucz line list.

The synthetic spectra are compared in Fig. 7. In general, the agreement is very good: the strengths and shapes of the line feature are well-matched, as is the total emergent intensity (in these test calculations, roughly 30 per cent of the radiation launched through the inner boundary is backscattered by the ejecta such that the total amplitude of the emergent spectrum is not trivially reproduced by imposing the same radiative flux at the inner boundary).

There are, however, some discrepancies. For example, the strength of the S II feature ($\sim$ 5300 Å) is noticeably different and there are small wavelength offsets: for example, Si II λ 6355, 5050 and the Ca II H+K / IR triplet features have larger blueshifts in the TARDIS spectrum. Many of these discrepancies can be attributed to subtle differences in the treatment of ionization. In particular, the modified nebular approximation used by PYTHON does not incorporate the correction factor δ proposed by ML93. If we do not include this correction (i.e. we force $\delta = 1$ in Equation 3), we find an even closer match between the two codes (see Fig. 7). Although still imperfect, this level of agreement between two independent codes gives confidence that both are performing well. Our comparisons also confirm that, although not qualitatively critical, even relatively minor changes to the ionization approximation are relevant to high-precision modelling.

5.4 Comparison to ARTIS

The ARTIS code is a purpose-build, multi-dimensional SN radiative transfer code. In contrast to TARDIS and PYTHON, ARTIS carries out time-dependent calculations that produce time sequences of synthetic spectra for an input model. ARTIS calculations do not involve an inner boundary through which radiation is injected – instead, the energy injection is followed in detail by simulating the emission of γ-rays origination from radioactive decays (see Lucy 2005; Sim 2007; Kromer & Sim 2009). Avoiding the need for an artificial inner boundary is the greatest advantage of codes such as ARTIS, but this comes at considerable computational expense: a single ARTIS calculation typically takes 10 – 100 kCPU hours, which is prohibitive for studies in which large parameter spaces are to be searched.

Like PYTHON, ARTIS includes bound-bound, bound-free, free-free and electron scattering opacities. ARTIS also includes an implementation of the macro atom scheme for all ions and an ionization approximation that is based on detailed photo-ionization rate estimators (see Kromer & Sim 2009, for details). As in TARDIS, the ejecta are assumed to be in homologous expansion.

For our comparisons between TARDIS and ARTIS, we use the same detonation model described in Section 4. The ARTIS simulations for the model yield a sequence of synthetic spectra covering the ultraviolet to infrared wavelength regions. We shall focus on comparisons at two epochs, 11.1 and 18 days after explosion (i.e. roughly at, and one week before, the time of maximum light in the optical bands). For the TARDIS comparison calculations, we use the luminosities calculated by ARTIS ($\log_{10} L_{\rm o}/_{\odot} = 9.34, 9.50$ for $t = 11.1, 18.0$ days, respectively) and adopt $v_{\rm o} = 22000$ km s^{-1}. We must also choose a location for the inner boundary of the TARDIS computational domain. Based on the model structure (Fig. 1), it seems inappropriate to consider $v_i \lesssim 8000$ km s^{-1}, since this would place the inner boundary within the ^{56}Ni-rich layers. Consequently, we have carried out several TARDIS runs for both epochs that explore the effect of choosing a range of values for $v_i \geq 9000$ km s^{-1}.

Synthetic optical and ultraviolet spectra from our TARDIS calculations are compared to ARTIS spectra in Fig. 8. First, we note that the choice of v_i does not significantly affect the optical spectrum, even when varied over a fairly wide range (9000 – 11000 km s^{-1}). This is reassuring since significant sensitivity to the choice of v_i would pose a challenge for the rapid exploration of model parameter space for which TARDIS is intended. If v_i is made very large ($v_i \gtrsim 12000$ km s^{-1}) consequences do appear, primarily in the ultraviolet and blue regions of the spectra, due to the lack of line-blocking by iron-group elements when only the high velocity ejecta are included (see Fig. 1).

Second, for a reasonable choice of v_i, the shapes and strengths of the optical spectral features agree fairly well in the two codes, suggesting that the simplifications made in TARDIS do not severely limit its ability to model optical lines, particularly for early epochs. Our comparison do, however, highlight important differences: in the TARDIS calculations there is more ultraviolet emission (blueward of $\sim$ 3000 Å) and correspondingly less optical emission. This discrepancy is significantly stronger at the later epoch considered – at 18 days, the optical pseudo continuum in ARTIS is clearly brighter than predicted by TARDIS. The origin of these differences is likely a combination of factors arising from the different approximations used. In Fig. 8, we show that both the radiation temperature ($T_{\rm R}$) and the mean intensity(described in terms of $T_J = W^{1/4} T_r$; see Kromer & Sim 2009) and fairly similar in both calculations.

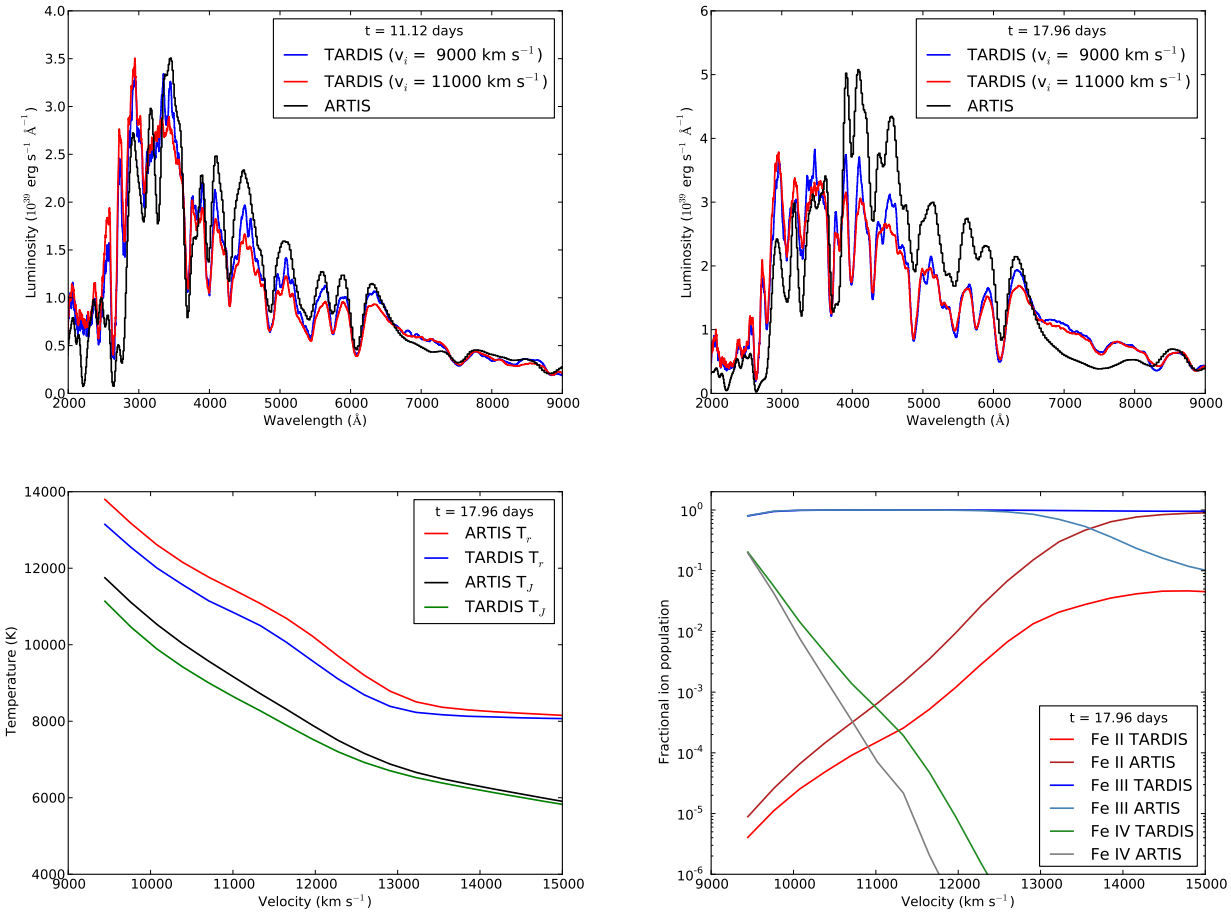

Figure 8. Comparisons of ARTIS and TARDIS calculations for the WD detonation model (see Section 4). The upper panels show synthetic spectra for 11.1 (left) and 18.0 (right) days after explosion. For TARDIS, spectra are shown for two calculations with different choices of $v_{\rm i}$. For the later epoch considered (18 days after explosion), the lower panels compare radiation field temperatures (left) and Fe ionization fractions (right) between ARTIS and the TARDIS calculation with $v_i = 9000$ km s^{-1}.

However, there is a distinct offset in the computed ionization state. For example, ionization fractions for Fe are shown in Fig. 8: while the overall pattern of the ionization is similar in both calculations, the ionization approximation used by ARTIS leads to systematically lower mean ionization. Since higher abundances of the singly ionized iron-group elements will lead to more effective line-blocking at blue wavelengths, this likely is responsible for much of the difference between the ultraviolet flux level in the two calculations.

Clearly, further study and improvement of the ionization approximations used is needed for quantitative modelling, and the limitations of these approximations must be borne in mind when comparing synthetic spectra from any of the codes discussed here to observations. Nevertheless, we conclude that, despite a considerable reduction in complexity, the TARDIS code is already capable of providing a reasonable approximation to the ARTIS calculations (for suitably chosen luminosity), lending credence to its use for initial attempts to fit spectra and explore model parameter spaces.

6 RESULTS

Thanks to its modularity, TARDIS is well suited to study the influences of different physical assumptions on the spectrum. Here we present results of simple calculations that investigate the differential effects of two important approximations that are sometimes adopted in the modelling of SN Ia spectra. For these tests we again use the sub-Chandrasekhar mass detonation model introduced in Section 4. We will show results computed for two epochs for which we adopt $L_{\rm o}$ and $v_{\rm i}$ as given in Table 3.

Table 3. Parameters used for the calculations in our comparisons between different physics modules. The luminosities adopted are extracted from the ARTIS simulations.

$t_{\rm exp}$ days	Luminosity $\log L_{\rm o}/L_{\odot}$	$v_{\rm i}$ km s^{-1}	$v_{\rm o}$ km s^{-1}
11.1	9.34	11000	22000
18.0	9.50	9000	22000

6.1 NLTE excitation

As noted in Section 1, several of the existing MC radiative transfer codes (including ML93, ARTIS and PYTHON) make use of simple approximations for the treatment of excited level populations. This has advantages for computational expediency and, particularly for

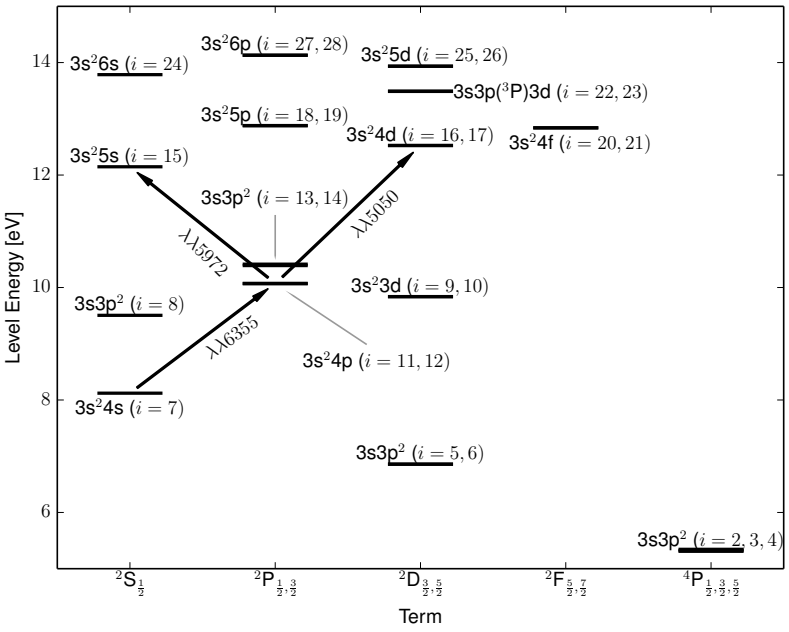

Figure 9. A Grotrian diagram of the Si II atomic model in the CHIANTI database.

multi-dimensional simulations, for reducing memory requirements but it comes at the cost of reduced accuracy. In contrast, studies by Pauldrach et al. (1996), Baron et al. (1996) and Blondin et al. (2013) have made use of the full equations of statistical equilibrium incorporating both ionization and excitation and shown that NLTE effects can be important. Here, we focus on exploring approximations for excitation (rather than ionization) with a view to quantifying the systematic errors introduced by the simple (and computationally cheap) `dilute-lte` formula to estimate excitation states.

We focus on four important ions, namely Si II, S II, Mg II and Ca II. Each of these exhibit prominent features in the spectra of SNe Ia and understanding potential systematic uncertainties in their modelling is an important ingredient in attempts to make synthetic fits to observed data sets. For our comparisons, we adopted the `nebular` ionization, `macroatom` line interaction and `detailed` radiative rates modes (i.e. the most sophisticated set of assumptions currently implemented) and then compared results obtained with `dilute-lte` and `nlte` excitation modes. We considered each of our selected ions in turn: i.e. for both epochs considered, we have computed a total of five synthetic spectra – one in which each of the four selected ions is treated in `nlte` excitation mode (all other ions in `dilute-lte` excitation mode) and one comparison calculation in which all ions are handled in `dilute-lte` excitation mode. This approach is convenient to assess the direct consequences of the level populations in key ions but we note that it does not capture any complex NLTE effects where one ion can influence the radiation field and therefore the excitation state of other ions. In all calculations, 10^7 MC quanta were used and 30 iterations were carried out. Synthetic spectra from these calculations are compared in Fig. 10.

To help quantify the effects of our different excitation treatment, we define a departure coefficient for multiplets

$$b^m_{\rm LTE} = \frac{\sum_i n_i/n_0}{\sum_i n_i^{\rm LTE}/n_0^{\rm LTE}}, \tag{23}$$

where i runs over the J sub-states of the term.

6.1.1 Silicon

The Si II $\lambda6355$ feature ($3s^2 4s\ ^2S \rightarrow 3s^2 4p\ ^2P$; see Fig. 9) is characteristic of SN Ia spectra and one of the main identifiers for this class. We find that the population of the $3s^2 4s\ ^2S$ state is relatively unaffected by our choice of excitation mode, resulting in a very small difference of departure coefficients between the two excitation treatments (see Table 4) and little change in the strength or shape of the spectral feature (Fig. 10).

In contrast, both the $\lambda5050$ ($3s^2 4p\ ^2P \rightarrow 3s^2 4d\ ^2D$) and the $\lambda5972$ ($3s^2 4p\ ^2P \rightarrow 3s^2 5s\ ^2S$) features are affected by the excitation mode: departure coefficients for $3s^2 4p\ ^2P$ are illustrated in Fig. 11 (and also listed in Table 4). Clearly our `nlte` excitation mode leads to significantly reduced populations for this state and correspondingly weaker spectral features compared to the `dilute-lte` approximation. E.g., in our 11.1 d (18.0 d) spectra, the equivalent width (EW) of $\lambda5972$ drops from 48 (38) to 24 (18) Å.

The different sensitivities of the departure coefficients for $3s^2 4s\ ^2S$ and $3s^2 4p\ ^2P$ mean that the choice of excitation treatment will systematically affect quantitative modelling of the equivalent width ratio $EW(\lambda6355)/EW(\lambda5972)$, which has been shown to act as a luminosity indicator (Nugent et al. 1995; Hachinger et al. 2008). Thus our TARDIS calculations verify that, as in Blondin et al. (2013), NLTE effects must be included if quantitative analysis (fac-

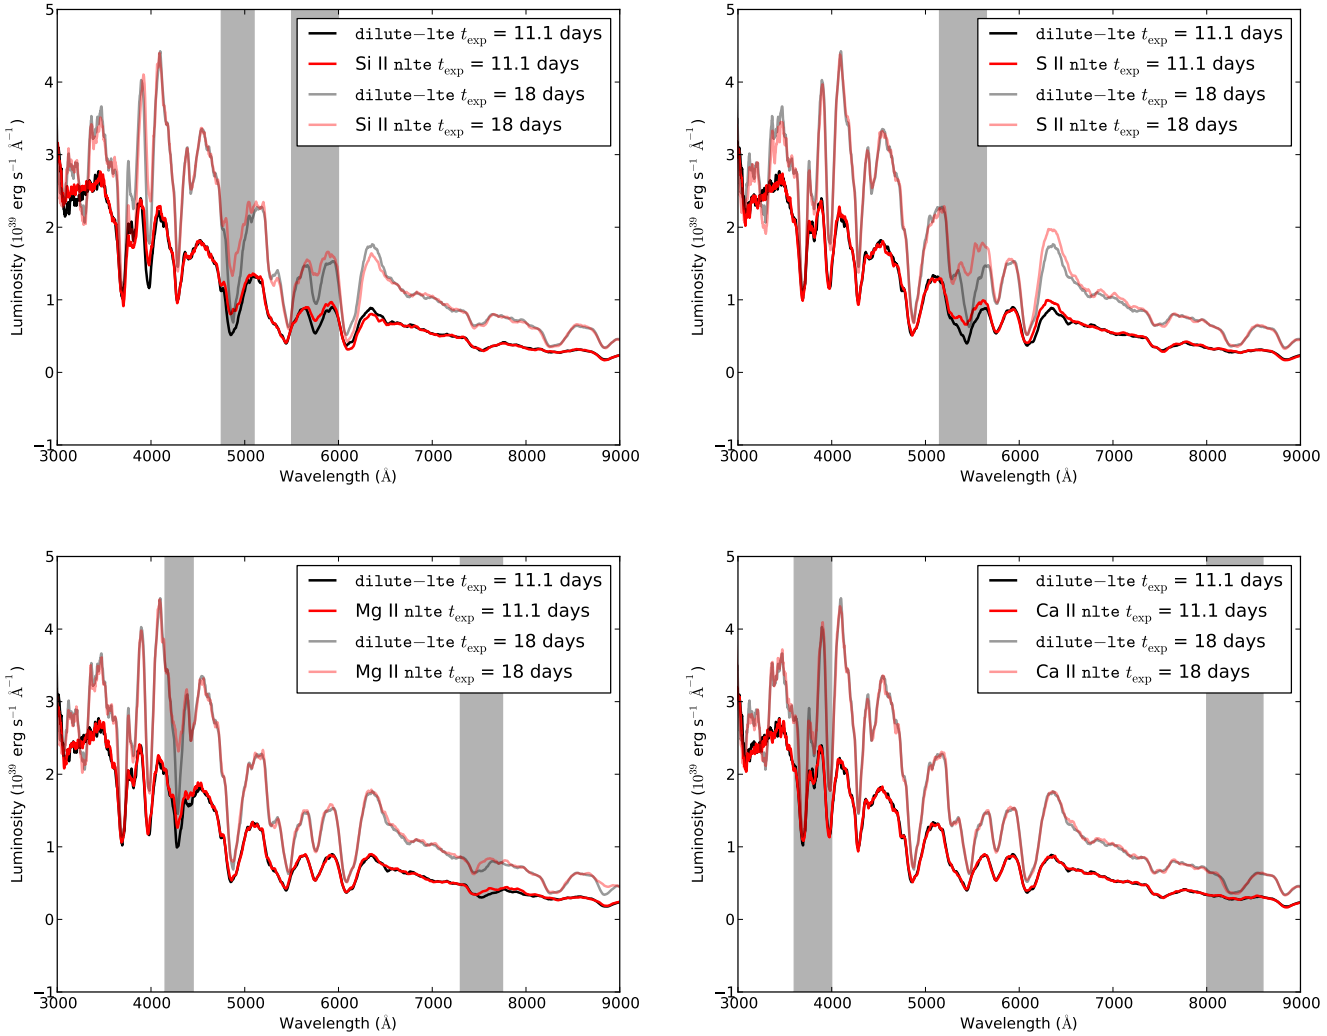
dilute−lte t_{exp} = 11.1 days
Si II nlte t_{exp} = 11.1 days
dilute−lte t_{exp} = 18 days
Si II nlte t_{exp} = 18 days
S II nlte t_{exp} = 11.1 days
S II nlte t_{exp} = 18 days
Mg II nlte t_{exp} = 11.1 days
Mg II nlte t_{exp} = 18 days
Ca II nlte t_{exp} = 11.1 days
Ca II nlte t_{exp} = 18 days
Luminosity (10^{39} erg s^{-1} Å^{-1})
Wavelength (Å)

Figure 10. Comparison of spectra computed using `dilute-lte` (black) and `nlte` (red) excitation modes for select species (see text). In each plot the regions most strongly affected are marked by a grey shade.

Table 4. Departure Coefficients

			Si II		S II		Ca II	Mg II	
t_{exp} day	shell velocity km s^{-1}	excitation treatment	4p ^{2}P	4s ^{4}P	3d ^{4}F	4s ^{2}S	3d ^{2}D	3d ^{2}D	4p ^{2}P
11.1	11000	`nlte`	0.55	0.35	0.42	0.40	0.83	0.25	0.21
		`dilute-lte`[a]	0.56	0.56	0.56	1.0	1.0	0.56	0.56
	17300	`nlte`	0.16	0.03	0.20	0.15	1.11	0.07	0.02
		`dilute-lte`[a]	0.19	0.19	0.19	1.0	1.0	0.19	0.19
18.0	9000	`nlte`	0.52	0.29	0.38	0.36	0.88	0.25	0.17
		`dilute-lte`[a]	0.49	0.49	0.49	1.00	1.0	0.49	0.49
	15500	`nlte`	0.20	0.05	0.24	0.21	1.74	0.02	0.01
		`dilute-lte`[a]	0.23	0.23	0.23	1.0	1.0	0.23	0.23

[a] The departure coefficient in `dilute-lte` model is 1 for meta-stable levels and the dilution factor W for other levels.

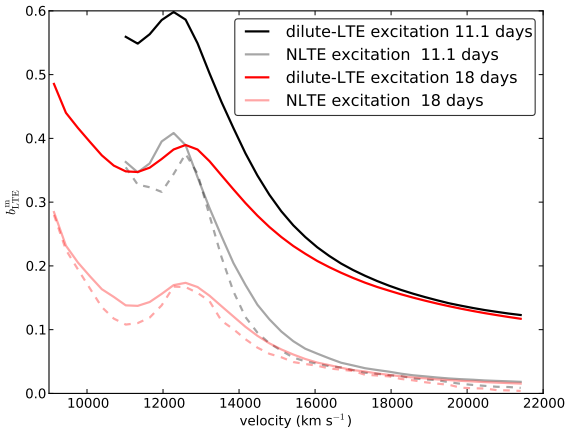

Figure 11. Comparison between the departure coefficients for Si II 4p ^{2}P using `dilute-lte` and `nlte` excitation modes. For `nlte` mode we show results using both `dilute-blackbody` (solid lines) and `detailed` (dashed lines) treatments of the radiative rates.

tor $\sim$ 2 accuracy or better) of this line ratio is to be used to test models.

6.1.2 *Sulphur*

S II is identified by iconic features around λ5449 and λ5623 in the spectra of SNe Ia (see Fig. 10): these are dominated by the $3s^2 3p^2(3P)4s\ ^4P \rightarrow 3s^2 3p^2(3P)4p\ ^4D$ and $3s^2 3p^2(3P)3d\ ^4F \rightarrow 3s^2 3p^2(3P)4p\ ^4D$ multiplets. As for the $3s^2 4p\ ^2P$ states in Si II, we find that our `nlte` excitation mode results in smaller level populations than `dilute-lte`, particularly for 3d ^{4}F, which is classified as meta-stable in the simpler scheme. This has quite noticeable ramifications for the shape of the S II feature (Fig. 10), again highlighting the sensitivity to excitation treatment in high-precision modelling.

6.1.3 *Magnesium*

Two particular features associated with Mg II could be of note in the modelling of optical SN Ia spectra: the strong λ4481 line (3d ^{2}D $\rightarrow$ 4f ^{2}F) and the 4p ^{2}P $\rightarrow$ 4d ^{2}D (λ7896) transition, which is a potential contaminant of the oxygen feature at λ7774. Once again we find that our `nlte` excitation treatment noticeably depopulates the lower levels of these transitions (compared to `dilute-lte`). This causes the λ4481 to become weaker (particularly at the later epoch, when it is less saturated) and the contribution of Mg II to the blend with λ7774 to nearly vanish. This lends confidence to the use of λ7774 as an oxygen abundance tracer (Blondin et al. 2013).

6.1.4 *Calcium*

Ca II imprints distinct spectral features at both ends of the optical spectrum: in the blue, the H and K lines (λ3950 arising from 4s ^{2}S $\rightarrow$ 4p ^{2}P) and in the near-infrared at λ8498, 8542, 8662 (3d ^{2}D $\rightarrow$ 4p ^{2}P). In both epochs considered, we find only a modest difference between departure coefficients in `nlte` and `dilute-lte` excitation modes for the 3d ^{2}D state (see Table 4), but this does not lead to a significant change in the shape or strength of the near-infrared feature (the absorption is already saturated). Thus, of those considered, the Ca II features are the least sensitive to the choice of excitation mode and are well-represented by the simple excitation formula.

6.2 Line interaction schemes

As described in Section 2.5, TARDIS allows for three different bound-bound interaction schemes: `scatter`, `downbranch` and `macroatom` modes. In this section, we explore the influence of the interaction scheme on the spectrum.

Mazzali (2000) showed that the ultraviolet flux is significantly affected if line branching is taken into account. When compared to calculations that include only resonance scattering, inverse fluorescence by iron group elements leads to enhanced emission in the ultraviolet. Our calculations also show this behaviour quite clearly – spectra calculated using `scatter` and `downbranch` modes are compared in Fig. 12.

With TARDIS we can also test the extent to which the macro atom scheme (see Section 2.5) alters the spectrum. In our calculations, the difference between `macroatom` (here combined with `detailed` mode for radiative rates) and `downbranch` mode is most apparent near maximum light (see Fig. 13) in the region around the S II features, $\sim$5500Å. The differences arise from the more complex pattern of frequency redistribution afforded by the macro atom scheme – this is illustrated in Fig. 13 where we show the increasing complexity of the frequency redistribution associated with the last line-interaction events that occurred for MC quanta in simulations with each of our line interaction modes.

To quantify the effects of the different line interaction methods we have measured equivalent widths of the S II feature (λ5449 and λ5623) and the Si II λ5987 line (see Table 5). We find that the former is sensitive to the line interaction method on the level of $\sim$ 30 Å but the later by only $\sim$ 3 Å.

6.3 Radiative rates estimation

Amongst the options currently implemented in TARDIS, the treatment of the bound-bound radiative rates has the most significant implications for the computational cost of a calculation. In `detailed` mode, it is necessary to have a sufficiently large number of MC quanta (many millions) to provide adequate statistics for the individual J^b_{lu} estimators (see Section 2.7) in every grid cell. Although more accurate, this is considerably more demanding than `dilute-blackbody` mode (Section 2.3) for which good convergence can be obtained with numbers of packets that are one to two orders of magnitude smaller. In this section we compare results from our two modes for bound-bound radiative rates to investigate whether the extra computation cost of `detailed` mode is warranted.

For the comparison, we repeat the Si II calculation from Section 6.1 adopting `dilute-blackbody` mode for the radiative rates. In Fig. 14, we compare values of J^b_{lu} for a selection of line transitions between the calculations with `detailed` and `dilute-blackbody` mode. We find that, although rather simplistic, the `dilute-blackbody` assumption for J^b_{lu} performs rather well – although there are modest deviations, the overall shape of the `detailed` calculation is well-matched, particularly in the inner regions where most of the spectrum formation occurs.

Given that J^b_{lu} is well-represented by the `dilute-blackbody` assumption, it is to expected that the

Table 5. Equivalent widths for spectral features measures from spectra calculated with various line interaction schemes

	Si II λ5987			S II λ5449 + λ5623		
$t_{\rm exp}$ days	`scatter` Å	`downbranch` Å	`macroatom` Å	`scatter` Å	`downbranch` Å	`macroatom` Å
11	41	37	43	153	142	171
18	50	38	39	158	143	175

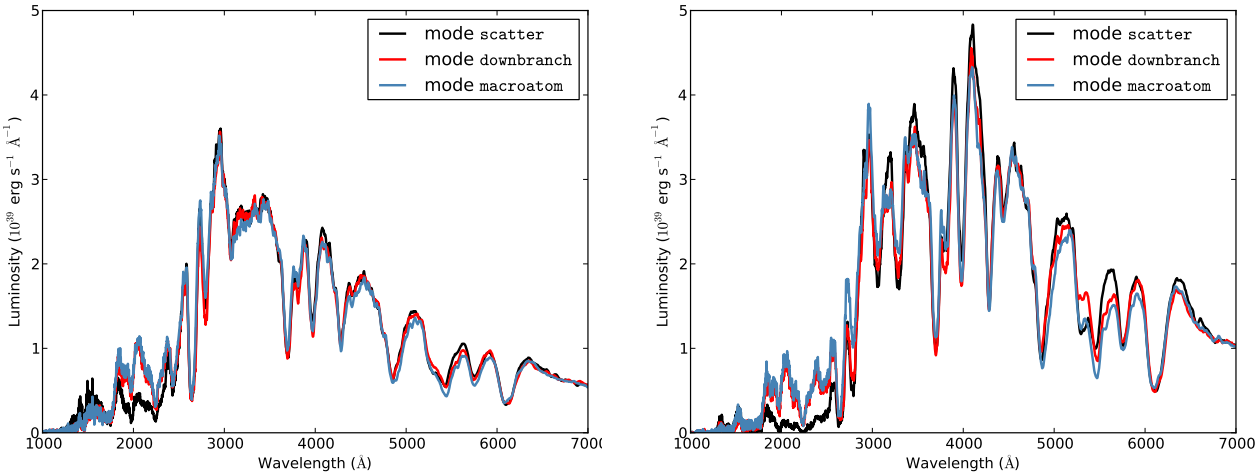

Figure 12. Comparison of TARDIS spectra computed for the detonation model with different line interaction schemes (`scatter`, `downbranch` and `macroatom` modes). The left panel shows results for $t_{\rm exp}$= 11.2 days (with $v_{\rm i}$ =11000 km s^{-1}) and the right panel for $t_{\rm exp}$= 18 days (with $v_{\rm i}$ = 9000 km s^{-1})

important quantities that depend on the bound-bound radiative rates will not be very adversely affected by this assumption. This is confirmed in Fig. 11, which compares departure coefficients for one example, and Fig. 15, in which the complete synthetic spectra computed using our two treatments of the radiative rates are shown. Consequently, we conclude that the `dilute-blackbody` assumption is generally acceptable and is therefore recommended for rapid modelling owing the considerable reduction in computational cost that it brings.

7 CONCLUSION & FUTURE WORK

We have presented a new 1D radiative MC code (TARDIS) for modelling of SN spectra, which is based on the indivisible-packet methods developed by Lucy (Abbott & Lucy 1985; Lucy & Abbott 1993; ML93; Lucy 1999, 2002, 2003). The purpose of the code is to allow rapid but accurate synthesis of SN spectra with only a few input parameters (see Appendix A). It is built in a modular fashion, making it possible for the user to activate or deactivate different physics and approximations so that it can be tuned for use in different applications.

Our successful comparisons with other codes (SYN++, PYTHON, ARTIS) verify that TARDIS is operating as expected and can accurately reproduce the shapes and strengths of complex spectral features obtained from independent calculations. Our results also highlight the importance of different ionization treatments in reconciling the results obtained with different codes and the ongoing need to pursue a good balance between accuracy and computational expediency in the treatment of ionization and excitation.

The modularity built into TARDIS allows for the exploration of different treatments of microphysics. In this work, we first explored the effect of simple excitation approximations (applied to four ions of relevance to SNe Ia, namely Si II, S II, Mg II, Ca II). We conclude that, for all of these except Ca II, the excitation treatment has a significant effect on the strength of line features which should be considered when attempting quantitative modelling of observations. We defer to a future study an equivalent investigation of approximate NLTE ionization treatments, which are likely to have more significant ramifications for the overall spectral shape (see Pauldrach et al. 1996; Baron et al. 1996; Blondin et al. 2013).

Secondly, we investigated the effect of different line interaction schemes on synthetic spectra. Similar to previous studies (Lucy 1999; Mazzali 2000) we found that the assumption of pure resonance scattering underestimates the ultraviolet luminosity of SNe Ia. Comparing the full macro atom approach to a simple down branching scheme shows that there is a minor but noticeable difference. As the computational overhead of using the macro atom scheme is only 10% above the simpler down branching, we advocate for the use of the macro atom scheme in future studies.

We also explored the sensitivity of TARDIS results to the different schemes for obtaining radiative rates for bound-bound transitions. Comparing results obtained using radiative rates recorded with line-by-line MC estimators to a simplified dilute-blackbody model (as adopted for the radiative rates in ARTIS), we see only minute changes in the synthetic spectrum. Given that the detailed estimators lead to an increase in computational time by orders of magnitude, our findings support the use of the simpler treatment.

TARDIS is now publicly available for use. For approximate modelling of SN Ia observations, we would generally recommend

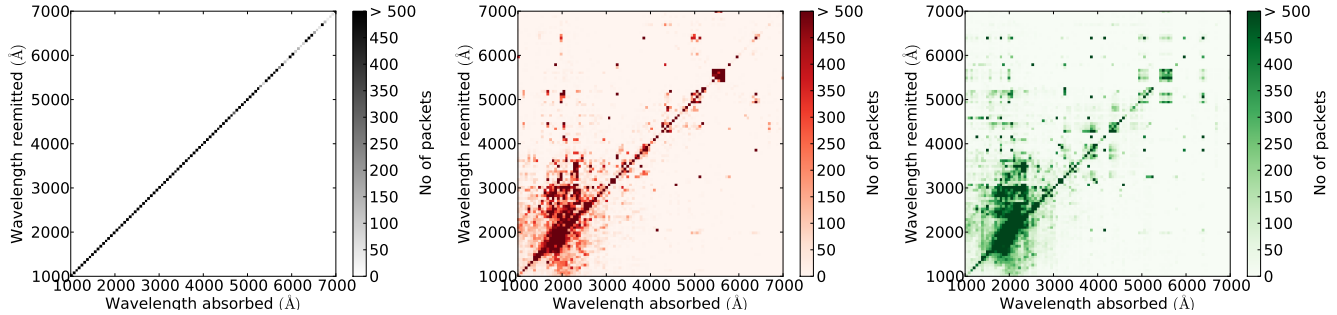

Figure 13. Each panel illustrates the wavelength redistribution that occurred during the last interaction of escaping packets (co-moving frame wavelength prior to interaction versus wavelength after interaction) for calculations at $t_{\rm exp} = 18$ days (`scatter` mode left, `downbranch` mode centre, `macroatom` mode right).

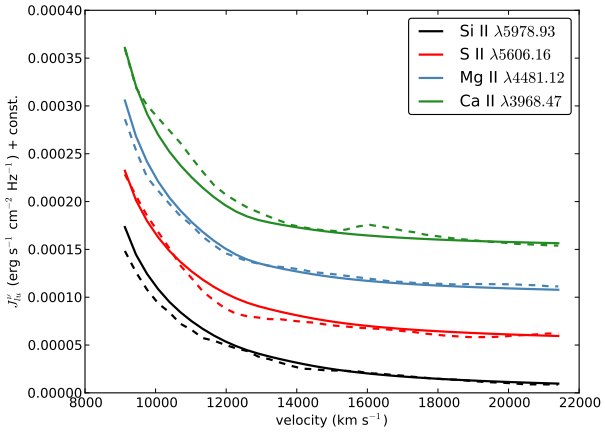

Figure 14. Comparison between J^b_{lu} obtained using the estimators in Equation 17 (dashed) and $J^b_{lu} = WB(\nu)$ (solid) for selected transitions at $t_{\rm exp}$=18 days. For clarity we have applied offsets to each curve (increment 5×10^{-5}).

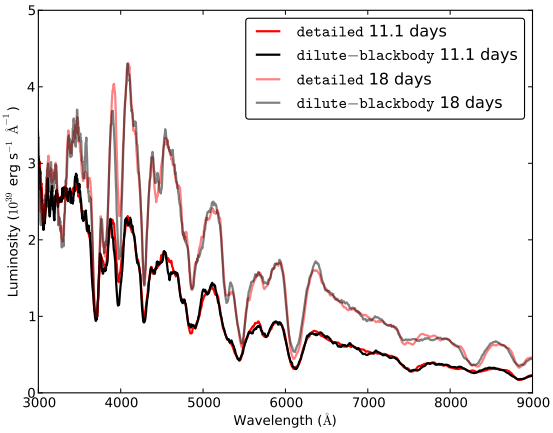

Figure 15. Comparison of spectra obtained for `detailed` and `dilute-blackbody` radiative rates modes (Si II treated in `nlte` excitatrion mode; all other ions in `dilute-lte` mode).

combining the `nebular` ionization / `dilute-lte` excitation plasma modes with the `dilute-blackbody` and `macroatom` interaction modes as a viable (computationally manageable) mode of operation. To fit a particular observation, a user would need to supply the luminosity and then develop a model by choosing a density profile [which could be empirical or based on an explosion model such as W7 (Nomoto et al. 1984)] and a set of ejecta abundances (which can be uniform or stratified). The density and abundances can then be modified to attempt to improve the fit (and therefore constrain the SN properties, as in e.g. Stehle et al. 2005; Mazzali et al. 2008; Hachinger et al. 2009; Tanaka et al. 2011). However, we stress that this is only one potential use/mode of operation for TARDIS and we encourage potential users to refer to the manual `http://tardis.rtfd.org` for further details.

In the near future, we plan to focus on two distinct TARDIS projects. First and foremost, the implementation of additional physics (bound-free/thermalization processes) with the goal of adding modules that include more sophisticated ionization approximations and allow for spectral synthesis for SNe II (Klauser & Kromer et al., in prep.). Secondly, since TARDIS was mainly developed to provide a means to fit SNe Ia with an approach similar to Mazzali et al. (2007), we aim to couple TARDIS with a suitable algorithm for automatic fitting of observations. We have explored this problem already, using the ML93 spectral synthesis code and genetic algorithms as the optimization algorithm and find that this is a promising approach (see Kerzendorf 2011, priv. comm. S. Hachinger, P. Mazzali).

8 ACKNOWLEDGEMENTS

We gratefully acknowledge Markus Kromer and Michael Klauser for many useful discussions and suggestions during all stages in the development of TARDIS. We thank Paolo Mazzali and Stephan Hachinger for helpful discussions of ionization/excitation treatments and formulation of the automatic spectral fitting project from which this work arose. We thank Knox Long for making available the PYTHON code, the LK02 ζ-values and for advice in making our comparisons with PYTHON. We are grateful to Ken Dere for help with CHIANTIPY and to Brian Schmidt for suggestions and support throughout the development of this project. We thank Adam Suban-Loewen for help in developing the TARDIS user interface and Aoife Boyle for participating in the testing of the code. We also would like to thank Erik Bray for helping with the installation frameworks implemented in TARDIS. We acknowledge the referee R. C. Thomas for very helpful suggestions and critical assessment of the paper. This research made use of Astropy, a community-developed core

Python package for Astronomy (Astropy Collaboration et al. 2013), and of the RANDOMKIT[6] library.

REFERENCES

Abbott D. C., Lucy L. B., 1985, ApJ, 288, 679
Anderson E. et al., 1999, LAPACK Users' Guide, third edn. Society for Industrial and Applied Mathematics, Philadelphia, PA
Astropy Collaboration et al., 2013, A&A, 558, A33
Baron E., Hauschildt P. H., Nugent P., Branch D., 1996, MNRAS, 283, 297
Blondin S., Dessart L., Hillier D. J., Khokhlov A. M., 2013, MNRAS, 429, 2127
Branch D., Doggett J. B., Nomoto K., Thielemann F.-K., 1985, ApJ, 294, 619
Dere K. P., Landi E., Mason H. E., Monsignori Fossi B. C., Young P. R., 1997, A&AS, 125, 149
Dessart L., Hillier D. J., 2005, A&A, 437, 667
Hachinger S., Mazzali P. A., Tanaka M., Hillebrandt W., Benetti S., 2008, MNRAS, 389, 1087
Hachinger S., Mazzali P. A., Taubenberger S., Pakmor R., Hillebrandt W., 2009, MNRAS, 399, 1238
Higginbottom N., Knigge C., Long K. S., Sim S. A., Matthews J. H., 2013, MNRAS, 436, 1390
Höflich P., 2003, in Hubeny I., Mihalas D., Werner K., eds, Astronomical Society of the Pacific Conference Series Vol. 288, Stellar Atmosphere Modeling. p. 185
Kasen D., Thomas R. C., Nugent P., 2006, ApJ, 651, 366
Kerzendorf W. E., 2011, PhD thesis, Australian National University, Research School of Astronomy & Astrophysics
Kramida A., Ralchenko Y., Reader J., NIST ASD Team 2012, Atomic Spectra Database (v5.0), National Institute of Standards and Technology, Gaithersburg, MD.
Kromer M., Sim S. A., 2009, MNRAS, 398, 1809
Kurucz R., Bell B., 1995, Atomic Line Data Kurucz CD-ROM No. 23. Cambridge, Mass.: Smithsonian Astrophysical Observatory
Lamers H. J. G. L. M., Cassinelli J. P., 1999, Introduction to Stellar Winds
Landi E., Del Zanna G., Young P. R., Dere K. P., Mason H. E., 2012, ApJ, 744, 99
Lentz E. J., Baron E., Branch D., Hauschildt P. H., 2001, ApJ, 557, 266
Long K. S., Knigge C., 2002, ApJ, 579, 725 (LK02)
Lucy L. B., 1999, A&A, 345, 211
Lucy L. B., 2002, A&A, 384, 725
Lucy L. B., 2003, A&A, 403, 261
Lucy L. B., 2005, A&A, 429, 19
Lucy L. B., Abbott D. C., 1993, ApJ, 405, 738
Matsumoto M., Nishimura T., 1998, ACM Trans. Model. Comput. Simul., 8, 3
Mazzali P. A., 2000, A&A, 363, 705
Mazzali P. A., Lucy L. B., 1993, A&A, 279, 447 (ML93)
Mazzali P. A., Röpke F. K., Benetti S., Hillebrandt W., 2007, Science, 315, 825
Mazzali P. A., Sauer D. N., Pastorello A., Benetti S., Hillebrandt W., 2008, MNRAS, 386, 1897
Noebauer U. M., Sim S. A., Kromer M., Röpke F. K., Hillebrandt W., 2012, MNRAS, 425, 1430
Nomoto K., Thielemann F.-K., Yokoi K., 1984, ApJ, 286, 644
Nugent P., Phillips M., Baron E., Branch D., Hauschildt P., 1995, ApJ, 455, L147
Pauldrach A. W. A., Duschinger M., Mazzali P. A., Puls J., Lennon M., Miller D. L., 1996, A&A, 312, 525
Pinto P. A., Eastman R. G., 2000, ApJ, 530, 757
Röpke F. K., 2005, A&A, 432, 969
Sauer D. N., Hoffmann T. L., Pauldrach A. W. A., 2006, A&A, 459, 229
Savitzky A., Golay M. J. E., 1964, Analytical Chemistry, 36, 1627
Seitenzahl I. R. et al., 2013, MNRAS, 429, 1156
Sim S. A., 2007, MNRAS, 375, 154
Sim S. A., Drew J. E., Long K. S., 2005, MNRAS, 363, 615
Sim S. A., Mazzali P. A., 2008, MNRAS, 385, 1681
Sim S. A., Miller L., Long K. S., Turner T. J., Reeves J. N., 2010, MNRAS, 404, 1369
Sim S. A., Röpke F. K., Hillebrandt W., Kromer M., Pakmor R., Fink M., Ruiter A. J., Seitenzahl I. R., 2010, ApJ, 714, L52
Stehle M., Mazzali P. A., Benetti S., Hillebrandt W., 2005, MNRAS, 360, 1231
Tanaka M., Mazzali P. A., Stanishev V., Maurer I., Kerzendorf W. E., Nomoto K., 2011, MNRAS, 410, 1725
Thomas R. C., Nugent P. E., Meza J. C., 2011, PASP, 123, 237
Wieser M. E., Coplen T. B., 2011, Pure and Applied Chemistry, 83, 359

APPENDIX A: USING TARDIS

At runtime TARDIS requires (1) an atomic database and (2) a user specified configuration file.

The atomic database is supplied in HDF5[7] format and consists of tables containing the data described in Section 3. A relatively simple atomic database is downloadable from the online manual (`http://tardis.rtfd.org` - in the "Running TARDIS" section). For custom versions of the atomic database users are currently encouraged to contact the authors.

The configuration file is supplied in YAML[8] markup language and defines both the physical and numerical parameters of the calculation and selects which of the available modes of operation are to be used. A simple example of the input file is shown and explained in Section A1 (further details in the code manual - `http://tardis.rtfd.org`).

In the standard mode of operation, the major physical parameters that can be varied are the output luminosity, the time since explosion and the density/composition profile of the model (including the locations of the inner and outer boundaries). The density profile can be chosen from preset standard options [e.g. constant density or a W7-like (Nomoto et al. 1984) density, as in the example given below] or it can be supplied via an addition ASCII input table (density tabulated as a function of velocity; see manual). The abundance distribution can be uniform (see example given below) or stratified (again using an ASCII input table; see manual).

[6] `http://js2007.free.fr/code/index.html#RandomKit`

[7] `http://www.hdfgroup.org/HDF5/`

[8] `http://www.yaml.org`

```
1  tardis_config_version: v1.0
2
3  supernova:
4      luminosity_requested: 2.8e9 Lsun
5      time_explosion: 13 day
6
7  atom_data: kurucz_atom_pure_simple.h5
8
9  model:
10
11     structure:
12         type: specific
13
14         velocity:
15             start: 1.1e4 km/s
16             stop: 2.0e4 km/s
17             num: 20
18
19         density:
20             type: branch85_w7
21
22     abundances:
23         type: uniform
24         O: 0.19
25         Mg: 0.03
26         Si: 0.52
27         S: 0.19
28         Ar: 0.04
29         Ca: 0.03
30
31 plasma:
32     ionization: nebular
33     excitation: dilute-lte
34     radiative_rates_type: dilute-blackbody
35     line_interaction_type: scatter
36
37 montecarlo:
38     seed: 23111963
39     no_of_packets : 2.0e+5
40     iterations: 30
41     last_no_of_packets: 5.0e+5
42     no_of_virtual_packets: 5
43
44 spectrum:
45     start: 500 angstrom
46     stop: 20000 angstrom
47     num: 10000
```

Figure A1. YAML input file for TARDIS.

A1 Example input file

Fig. A1 shows a configuration file for TARDIS. This file initiates a calculation equivalent to that used in Section 5.3 (i.e. see Fig. 7) and runs successfully with version 0.9 of TARDIS. The sections of the file are described in detail below (with reference to line numbers). We stress that this is only one simple example and refer users to the manual (`http://tardis.rtfd.org`) for further information.

All dimensional quantities in the configuration file require a unit [these are parsed using the ASTROPY (Astropy Collaboration et al. 2013) `Quantity` framework].

The first line of the configuration file identifies the TARDIS version. Currently only `v1.0` is supported but future versions of TARDIS may involve a restructure of the configuration file and will recognise the layout according to the version number.

The section (`supernova`) defines the requested emergent luminosity ($L_{\rm o}$; see Section 2.8) and the time since explosion ($t_{\rm exp}$). The standard ASTROPY units are available.

The atomic data file (HDF5 format) is specified at line 7 (absolute or relative path can be given).

The `model` section controls the setup of the radial density and abundance profiles. In the example, we specify 20 shells (line 17) equally spaced in the velocity range 11000 km s^{-1} to 20000 km s^{-1} (lines 15–16). We adopt a W7-like density profile (specifically, a polynomial fit to the W7 structure similar to Branch et al. 1985; lines 19–20). The example calculation adopts uniform (line 23) composition of all zones by specifying mass fractions for six elements in lines 24–29. Abundances of elements not specified are set to zero. More complex radial stratification of density and/or composition can be specified (see user manual).

The settings to be used in determining the ionization/excitation/line-interactions modes are provided in the `plasma` section (lines 31–35) using the keywords described in Table 1.

In the `montecarlo` section, TARDIS is supplied with the seed value for the random number generator (line 38), the requested number of packets to be used in each MC simulation (line 39) and the number of iterations to be performed in determining the plasma state (line 40). In addition (line 41), one can specify a different number of packets for the last MC simulation (making it possible to increase the signal-to-noise of the final spectrum). If the keyword `no_of_virtual_packets` is greater than zero (line 42), the virtual packet mode will be activated in the last MC simulation with N_v set to the value given (see Section 2.9).

In the last section (`spectrum`), the spectral range for the synthetic spectrum is set by specifying start and end wavelengths (lines 45 and 46). The spectral resolution is determined by choosing the total number of wavelength bins (line 47).

Running the model defined by the example configuration file on a mid 2012 MacBook Air (Intel Core i7 2.00 GHz, 8 GB RAM) takes approximately $\sim$ 3 minutes to complete all 30 iterations. The runtime scales roughly linearly with the number of MC packets.